

1. In reconstruction imaging in which a volume is reconstructed from a series of measured projection images, each generated by projection of a radiation source positioned at a respective focus through the volume to a detector at which the respective measured projection image is acquired ("detector plane"), an improved method of back-projecting a two-dimensional (2D) representation ("first 2D representation") to generate three-dimensional (3D) representation, where the first 2D representation is in the detector plane, the improvement comprising

for each of one or more slices of the 3D representation parallel to the detector plane and for each distinct focus at which a said measured projection images is generated, performing the steps of:

- i) warping the first 2D representation to generate a second 2D representation, the warping step including applying to the first 2D representation a selected linear mapping, where that selected linear mapping would map, in order to match dimensions of the respective slice within the 3D representation, a region defined by projection, at the respective focus, of corners of that slice onto the detector plane,
 - ii) incrementing values of each of one or more voxels of the respective slice by an amount that is a function of a value of a correspondingly indexed pixel of the second 2D representation.
2. In the reconstruction imaging of claim 1, the further improvement comprising

forward-projecting the three-dimensional (3D) representation ("first 3D representation") to generate a two-dimensional (2D) hypothetical projection image, the forward-projecting step comprising

for each distinct x-ray source focus at which a respective measured projection image is generated and for each of one or more respective slices in the first 3D representation parallel to the detector plane, performing the steps of

- i) warping the respective slice of the first 3D representation to generate a respective slice of a second 3D representation, the warping step including applying to the respective slice of the first 3D representation a selected linear mapping, where

that selected linear mapping would map a region defined by projection of corners of the respective measured projection image onto the respective slice of the first 3D representation for the respective x-ray source focus to match dimensions of the respective measured projection image in the detector plane,

- ii) incrementing values of each of one or more pixels of the 2D representation by an amount that is a function of a value of a correspondingly indexed pixel of the respective slice of the second 3D representation.
- 3. In the reconstruction imaging of claim 1, the further improvement comprising executing at least the warping step on a graphics processing unit (GPU) coprocessor.
 - 4. In the reconstruction imaging of claim 3, the further improvement wherein the GPU is of the type having a programmable pixel shader and a programmable vertex shader.
 - 5. In the reconstruction imaging of claim 3, wherein the GPU has one or more instructions that provide for parallel processing of multi-component operands, the further improvement comprising executing such an instruction with a lower-order portion of a mantissa in one component of an operand and a higher-order portion of that mantissa in another component of that same operand.
 - 6. In the reconstruction imaging of claim 5, the further improvement wherein the instruction is an add instruction.
 - 7. In reconstruction imaging in which a volume is reconstructed from one or more projection images, each generated by projection of a radiation source positioned at a respective focus through the volume to a plane at which the respective projection image is acquired ("detector plane"), an improved method of forward-projecting a three-dimensional (3D) representation ("first 3D representation") to generate a two-dimensional (2D) representation, where the 2D representation is in the detector plane, the improvement comprising

for each distinct focus at which a said projection image is generated and for each of one or more slices in the first 3D representation parallel to the detector plane, performing the steps of

- i) warping the respective slice of the first 3D representation to generate a respective slice of second 3D representation, the warping step including applying to the respective slice of the 3D representation a selected linear mapping, where that selected linear mapping would map a region defined by projection of corners of the respective projection image onto the respective slice for the respective focus to match dimensions of the projection image in the detector plane,
 - ii) incrementing values of each of one or more pixels of the 2D representation by an amount that is a function of a value of a correspondingly indexed pixel of the respective slice of the second 3D representation.
- 8. In the reconstruction imaging of claim 7, the further improvement comprising executing at least the warping step on a graphics processing unit (GPU) coprocessor.
- 9. In the reconstruction imaging of claim 8, the further improvement wherein the GPU is of the type having a programmable pixel shader and a programmable vertex shader.
- 10. In the reconstruction imaging of claim 9 wherein the GPU has one or more instructions that provide for parallel processing of multi-component operands, the further improvement comprising executing such an instruction with a lower-order portion of a mantissa in one component of an operand and a higher-order portion of that mantissa in another component of that same operand, thereby allowing one or more operations to be performed at higher precision than the underlying computational hardware was otherwise capable of.
- 11. In the reconstruction imaging of claim 10, the further improvement wherein the instruction is an add instruction.
- 12. In a method of medical imaging in which a volume is reconstructed from a series of projection images, each generated by projection of a radiation source positioned at a respective focus through the volume to a detector plane at which the respective projection is acquired ("detector plane"), the improvement comprising the steps of:
 - A. for each focus, forward-projecting a three-dimensional (3D) representation ("first 3D representation" that is an estimate of the volume) to generate a two-dimensional (2D)

representation ("first 2D representation") that is an estimate of the projection acquired with the source positioned at that respective focus,

the forward-projecting step including warping one or more slices of the first 3D representation to generate corresponding slices of a second 3D representation by mapping voxels of the respective slice of the first 3D representation to voxels of the corresponding slice of the second 3D representation as if rays from the source, when positioned at the respective focus, projected through the volume at a substantially constant angle normal to the detector plane of the respective projection, where the 2D representation is in the detector plane and where the slices of the first and second 3D representation are parallel to the detector plane,

- B. for each of the selected focus, back-projecting a 2D representation ("second 2D representation") that is a function of the first 2D representation and a respective projection to generate a third 3D representation,

the back-projecting step including warping the second 2D representation to generate a third 2D representation by mapping pixels of the second 2D representation to pixels of the third 2D representation as if rays from the source, when positioned at the respective focus, projected through the volume at a substantially constant angle normal to the detector plane of the respective projection, where the second and third 2D representations are in the detector plane.

13. In the method of claim 12, the further improvement wherein the method is an iterative reconstruction method.
14. In the method of claim 13, the further improvement wherein the method as a maximum likelihood estimation maximization method.
15. In the method of claim 14, the further improvement wherein the second 2D representation is a difference between the first 2D representation and the respective projection.
16. In the method of claim 14, the further improvement comprising modifying values of the first 3D representation as a function of the third 3D representation.

17. In the method of claim 12, the further improvement wherein the warping step of the forward-projecting step includes applying to the respective slice of the first 3D representation a selected linear mapping, where that selected linear mapping would map a region defined by projection of corners of the respective projection onto the respective slice for the respective focus to match dimensions of the projection in the detector plane.
18. In the method of claim 17, the further improvement wherein the forward-projecting step includes incrementing values of each of one or more pixels of the first 2D representation by an amount that is a function of a value of a correspondingly indexed pixel of one or more respective selected slices of the second 3D representation.
19. In the method of claim 12, the further improvement wherein

the warping step of the back-projecting step includes applying to the applying to the second 2D representation a selected linear mapping, where that selected linear mapping would map a region defined by projection of corners of the respective slice of the third 3D representation onto the detector plane for the respective focus to match dimensions of the respective slice within the third 3D representation.
20. In the method of claim 19, the further improvement wherein the back-projecting step includes incrementing values of each of one or more voxels of each of at least selected slices of the third 3D representation by an amount that is a function of a value of a correspondingly indexed pixel of the second 2D representation.
21. In the method of claim 12, the further improvement comprising executing at least the warping steps on a graphics processing unit (GPU) coprocessor.
22. In the method of claim 21, the further improvement wherein the GPU is of the type having a programmable pixel shader and a programmable vertex shader.
23. In the method of claim 21, wherein the GPU has one or more instructions that provide for parallel processing of multi-component operands, the further improvement comprising executing such an instruction with a lower-order portion of a mantissa in one component of an operand and a higher-order portion of that mantissa in another component of that same operand.

24. In the method of claim 23, the further improvement wherein the instruction is an add instruction.